INVITED PAPER *Special Issue on Distributed Processing for Controlling Telecommunications Systems*

Comparison of Adaptive Internet Multimedia Applications*

Xin WANG† *and* **Henning SCHULZRINNE**†, ***Nonmembers***

SUMMARY The current Internet does not offer any quality of service guarantees or support to Internet multimedia applications such as Internet telephony and video-conferencing, due to the best-effort nature of the Internet. Their performance may be adversely affected by network congestion. Also, since these applications commonly employ the UDP transport protocol, which lacks congestion control mechanisms, they may severely overload the network and starve other applications. We present an overview of recent research efforts in developing adaptive delivery models for Internet multimedia applications, which dynamically adjust the transmission rate according to network conditions. We classify the approaches used to develop adaptive delivery models with brief descriptions of representative research work. We then evaluate the approaches based on important design issues and performance criteria, such as the scalability of the control mechanism, responsiveness in detecting and reacting to congestion, and ability to accommodate receiver heterogeneity. Some conclusions are developed regarding the suitability of particular design choices under various conditions.
key words: *adaptive, multimedia, internet, sender-driven, receiver-driven transcoder, QoS*

1. Introduction

The development and use of distributed multimedia applications are growing rapidly at present. Some common examples are video-conferencing, Internet telephony, and video-on-demand. High-quality delivery of multimedia information requires high network bandwidth. Also, since a minimum audio and video quality are required in order to communicate the desired information, video and audio applications require a certain minimum throughput for useful operation. Finally, in order to support interactive conversations, and to ensure synchronization of data belonging to different streams (for example, audio and video), as well as within a stream, there should be an upper bound on the end-to-end delay, and on the maximum variation in delay.

The special characteristics of multimedia applications place a number of requirements on the network. The requirements can be specified in terms of quality of service (QoS) parameters, such as throughput, packet loss, delay, and jitter. In a network providing undifferentiated, best-effort service without any QoS support mechanisms, fluctuations in network load can adversely affect multimedia applications. Also, multimedia applications on the Internet commonly employ the UDP transport protocol, which lacks any congestion control mechanism. As a result, applications with high bandwidth can severely overload the Internet, and starve TCP applications (which perform congestion control) of their fair share of bandwidth.

Different approaches may be considered to address these shortcomings. One approach is to enhance the network with mechanisms such as resource reservation [2] [3], admission control [4], and special scheduling mechanisms [5], so that a certain level of QoS can be guaranteed to an application. A certain degree of QoS support can also be provided by allowing differentiated or prioritized service at network switches [6].

Another approach is to adjust the bandwidth used by an application according to the existing network conditions. This approach has the advantage of better utilizing available network resources (which change with time), compared to approaches relying on resource reservation. It is also facilitated by the nature of existing multimedia applications, many of which allow the media rate and quality to be adjusted over a wide range. At the same time, the special requirements of multimedia applications mean that strictly TCP-like congestion control may not be suitable for these applications. The rate is halved (to a first order approximation) for every lost packet in TCP congestion control, which may cause corresponding sharp changes in encoder parameters to achieve the desired rate, and unpleasant perceived quality. Too small a rate may also violate the minimum throughput requirements of the application. Also, the per-packet acknowledgments used in TCP are not appropriate. The strict delay and delay jitter constraints may not allow lost packets to be re-transmitted, and multimedia applications in general can tolerate a small amount of packet loss. Per-packet acknowledgments also impose a large bandwidth overhead in a high-bandwidth multimedia application.

Approaches such as resource reservation and rate adaptation may be used together. In this paper, we restrict ourselves to an overview of recent research which focuses mainly on adaptive control schemes which regulate the rate of a multimedia application according to network conditions.

This review is not intended to be exhaustive. Our

Manuscript received October 27, 1998.
†The authors are with the Department of Computer Science, Columbia University, 1214 Amsterdam Avenue, New York, NY 10027, USA.
*Parts of this review are based on material from a tutorial at [1].

goal is to review work which is representative of ongoing research in this field, and evaluate the suitability of the approaches discussed under various conditions.

Adaptive control schemes presented in the literature can be broadly classified into sender-driven (Sect. 3), receiver-driven (Sect. 4), and transcoder-based (Sect. 5). Sender-driven schemes require the sender to respond to fluctuations in the service available from the network, and adjust its transmission accordingly. Receiver-driven schemes specify a mechanism for each receiver to select transmission of a particular quality according to the service it receives from the network. Transcoder based schemes place gateways at appropriate locations to deliver different levels of quality to network regions with different types of connectivity or different levels of congestion.

A number of other design alternatives and goals need to be considered in developing rate-adaptive control schemes. Some important issues are:

1. the signaling or feedback mechanism used to convey congestion information, which in turn drives the transmission rate adaptation process;
2. the specific rate control mechanism used in response to feedback;
3. the responsiveness of the congestion control scheme in detecting and reacting to network congestion;
4. the capability of the scheme to accommodate a diverse group of receivers that differ in their connectivity to the network, the amount of congestion on their delivery paths, and their need for transmission quality;
5. the scalability of the control mechanism in a multicast session with a large number of receivers;
6. fair sharing of bandwidth with competing connections, particularly TCP connections;
7. the perceived quality of received multimedia streams.

2. Sender-Driven Adaptation

Sender-driven adaptation schemes that are discussed here fall into two categories. Buffer-based adaption schemes use the occupancy of a buffer on the transmission path as a measure of congestion. Loss based adaptation schemes adjust the rate based on the packet loss experienced by receivers. Adaptation schemes have been proposed based on other congestion indicators, including CSMA/CD collisions [7], packet delay [7] and delay jitter [8]. Due to space constraints, we restrict our discussion to the more commonly used buffer occupancy based and loss based mechanisms.

2.1 Buffer-Based Adaptation

Buffer-based adaptation schemes base the adaptation of the transmission rate on the occupancy of a buffer on the transmission path. Essentially, the goal of the control algorithm is to maintain buffer occupancy at a constant, desired level. When the buffer begins to fill up, the transmission rate is reduced in response, and when the buffer begins to empty, the transmission rate is increased.

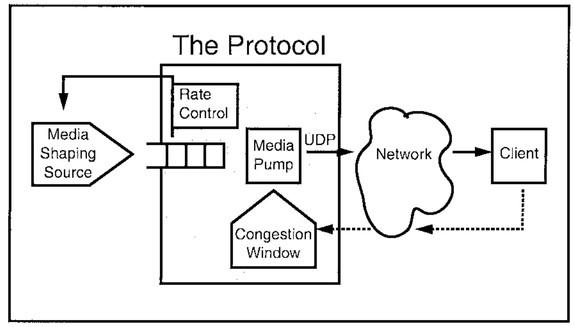

Fig. 1 Protocol framework used in buffer occupancy based adaptation scheme of Jacobs et al. [10].

Kanakia, Mishra and Reibman (KMA)[9] describe a scheme in which the sender periodically receives explicit feedback from the network giving the buffer occupancy and service rate received by the connection at the bottleneck queue. To account for the latency of the feedback, the evolution of the current bottleneck buffer occupancy and service rate are estimated. The estimates are used by a proportional control system to calculate the target transmission rate prior to transmitting each video frame. In order to meet the targeted sending rate, the quantization (Q) factor of the encoder is adjusted suitably. A damping mechanism is used to prevent sudden changes in the Q factor, and thus prevent annoyingly sudden changes in the perceived quality. If the transmission uses MPEG encoding, a separate service rate estimation is maintained for each type of frame (I, B, and P) by keeping a separate service rate estimator for each.

Jacobs and Eleftheriadis [10], [11] (JE) propose a protocol that uses the TCP congestion window (and hence, TCP acknowledgment messages from the receiver) to monitor congestion in the network. The stated goal of the authors is to allow video transmissions to adapt to network congestion in a manner similar to TCP, and thus ensure that the protocol competes fairly with TCP connections for available bandwidth. The TCP window size is used to govern the output rate of packets to the network. Packets are stored in a local buffer prior to sending, and the occupancy of this buffer in turn drives a proportional derivative control loop to determine the desired encoder rate (Fig. 1). *Dynamic Rate Shaping* [12] is used to adjust the encoder rate. The DRS operation reduces the source rate by eliminating a set of Discrete Cosine Transform (DCT) coefficients using a Lagrangian optimization.

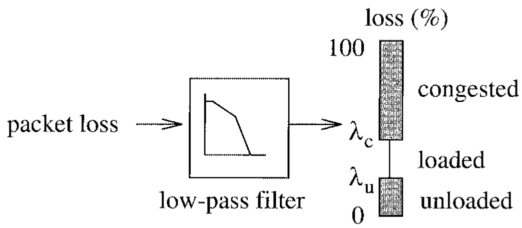

Fig. 2 Linear regulator with dead zone (from [15]).

2.2 Loss-Based Adaptation

Loss based adaptation schemes [15], [16], [21] regulate the transmission rate based on loss rate information reported by the receivers. The schemes discussed in this section consider the problem of controlling a video transmission multicast to a group of receivers over a packet switched network. Qualitatively, all three of the adaptation schemes adopt the following approach (Fig. 2). Based on feedback information from a receiver, the sender assumes that the receiver is in one of three states: *unloaded*, *loaded*, or *congested*. In the *unloaded* state, the sender progressively increases its transmission rate in an additive manner in response to feedback, until the network state is driven into the *loaded* state or the sender is sending the maximum useful rate. In the *loaded* state, the sender maintains a constant transmission rate. Depending on packet loss feedback, it can be driven into either the *unloaded* or the *congested state*. In the *congested* state, the sender progressively reduces its transmission rate multiplicatively until the reported loss decreases to the *loaded* state.

Issues that need to be considered include the loss thresholds for determining a particular network state, and the parameters controlling the additive rate increase and multiplicative rate decrease. In a multicast environment with heterogeneously connected receivers, different receivers may experience widely varying degrees of congestion. The sender must deal with the problem of deciding upon an overall network state based on feedback from these receivers. The adaptation schemes take different approaches in tackling this problem, and these are discussed in Sect. 5.4

In an adaptation scheme proposed by Bolot, Turletti, and Wakeman [16], the sender determines the network state perceived by receivers through a scalable feedback mechanism by using a probabilistic polling method that avoids the generation of feedback messages by every receiver in the multicast group. We return to a detailed discussion of the probabilistic polling scheme in Sect. 5.2. The authors specifically consider using the adaptation scheme to control the output rate of an H.261 encoder by adjusting the frame rate or the quantizer and movement detection threshold.

In algorithms proposed by Busse, Deffner and Schulzrinne [15], and Sisalem and Schulzrinne [21], the receiver reports of the RTP [14] control protocol (RTCP) are used to convey feedback to the sender, allowing the sender to calculate the packet loss and round trip delay for each receiver. In [15], a low-pass filter is used to smooth the reported packet loss rate, and the smoothed packet loss rate is used to determine the current congestion state, again using the model of Fig. 2. In the Loss-Delay based Algorithm (LDA) [21], in addition to the RTCP feedback, the sender estimates the bottleneck bandwidth of each receiver through a separate mechanism. Also, unlike the other two schemes, the additive rate is not fixed, and depends on the estimated bottleneck bandwidth, and the round-trip delay. For each receiver i, the sender uses the receiver feedback to compute a desired rate r_i, and an additive rate increase parameter AIR_i. Depending on the reported loss, each r_i is either increased additively or reduced multiplicatively, and the sender periodically searches the set of transmission rates r_i of all receivers and sets the output r to the minimum rate in the set.

Both of the above methods ([15] and [21]) adjust the frame rate at the encoder in order to achieve a specific output rate.

3. Receiver-Driven Adaptation

In receiver driven adaptation, receivers individually tune the received transmission according to their needs and capabilities. A number of receiver driven schemes use a combination of layered encoding, and a layered transmission scheme. The source data is encoded into a number of layers. A base layer provides the minimal QoS needed for an acceptable representation of the original data stream. Incrementally combining higher layers with the base layer results in a progressively higher QoS. Each encoded layer is transmitted to a separate multicast group. An alternative to this cumulative layering scheme is to encode and transmit multiple copies of the source input; each copy is encoded to have a different level of QoS, and sent to a separate multicast group. Although this approach (commonly referred to as simulcast) makes inefficient use of bandwidth, it may be more appropriate for layered transmission of audio using separate encoders for each layer, since audio encoders usually do not support layered encoding. Possibly due to the more demanding nature of distributed video applications relative to audio, layered adaptive schemes reported in literature generally use the former (cumulative layering) approach, and we restrict ourselves to this class of schemes in the present discussion.

The receiver selects a transmission quality appropriate to its requirements and constraints by subscribing to a certain number of multicast groups carrying different layers. The receiver monitors network congestion (based on parameters such as packet loss and throughput), and adapts to changes in network condi-

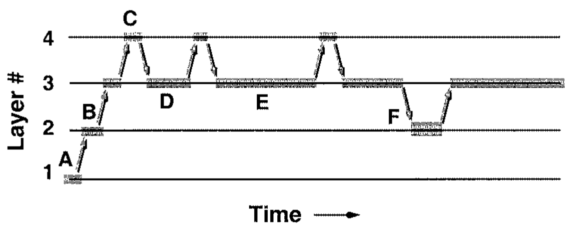

Fig. 3 A simple illustration of adaptive behavior in RLM [18]. Following an initial rapid series of joins, a join experiment to layer 4 fails at C. The exponential back-off of the join timer results in successively longer delays before join experiments (D and E).

tions by adding or dropping layers accordingly.

3.1 Receiver-Driven Adaptation without Rate Adjustment

In the RLM scheme proposed by McCanne, Jacobson, and Vetterli [18], the sender takes no active role in the adaptation mechanism. It encodes the source signal into cumulative layers, and transmits each layer of the signal to a separate IP multicast group. When the packet loss exceeds a certain threshold, the receiver perceives congestion and drops a layer. In an uncongested state, the receiver conducts *join-experiments* at intervals: when a *join timer* for the lowest unsubscribed layer goes off, the receiver subscribes to that layer. If this results in congestion within a *detection-time* period, the receiver reverts to its previous subscription level, and also multiplicatively increases the timer for the level associated with the failed join experiment. Otherwise, the receiver maintains the new level. When a receiver remains at a particular level without congestion, it multiplicatively reduces the associated join timer at intervals. This adaptive behavior is illustrated in Fig. 3.

The *ThinStreams* adaptation scheme [19], proposed by Wu, Sharma, and Smith, is similar to RLM in its overall framework, and we discuss mainly the important differences between the two. A primary distinction is that in the ThinStreams scheme, the granularity with which receivers may add/drop layers is decoupled from the granularity with which the source signal is encoded into layers. Each encoded layer at the sender is termed a thick stream, and this is split up into several *thin streams* of a fixed, small bandwidth. Each thin stream is sent to a separate multicast group. A receiver drops and adds thin stream layers based on perceived network congestion. The stated goal behind this refinement is that in a layered transmission scheme, the encoded layers can have large bandwidth. When a receiver adds a layer (for example, in a join experiment of RLM) it may overload a channel, resulting in significant lost packets for the receiver as well as other users sharing the bottleneck. Over time, this could also result in large oscillations in network congestion and in the quality perceived by receivers. By experimentally adding bandwidth in small fixed increments, the receiver prevents excessive overloading of the channel during join experiments. The ThinStreams algorithm uses the difference between the expected and measured throughput as the indicator of congestion. When the receiver is congested, it drops the group corresponding to the highest layer.

3.2 Receiver-Based Adaptation with Rate Adjustment

Layered encoding of video usually results in a small number of high bandwidth layers. Adaptation by adding or dropping an encoded layer is of a correspondingly large granularity, and this may result in underutilization of bandwidth, and sub-optimal quality of reception. One approach towards alleviating this problem is taken by the ThinStreams protocol. An alternative approach is to have the source dynamically adjust the bandwidth of each encoded layer in response to feedback from the receivers or the network. In this section we discuss adaptation schemes which combine the layered encoding and transmission architecture used in receiver based adaptation schemes such as RLM and ThinStreams, with rate adaptation by the sender in response to feedback.

Sisalem and Emanuel [22] propose an Adaptive Layered Transmission (ALT) protocol. The sender monitors loss information for each layer through periodic RTCP receiver reports. The transmission rate of each layer is adapted using the additive increase/multiplicative decrease model used by the loss based adaptation schemes discussed earlier (Fig. 2). In addition to the rate adaptation by the sender, if a receiver experiences packet loss above a certain threshold, it drops a layer to avoid driving the transmission rate of the layer down too low. If the receiver determines that it has excess capacity, it adds a layer. If all the receivers drop the current highest layer, or if the transmission rate of the highest layer is reduced below the minimum transmission rate, the sender may choose to temporarily discontinue the layer.

In [28], Vickers, Albuquerque, and Suda propose a rate-based adaptation schems, as well as a credit-based scheme (AMML). In response to receiver feedback, the sender decides the number of layers to encode, and the rate at which to transmit each layer. In both cases, the network is assumed to provide prioritized service. The base video layer has the highest priority, and successive enhancement layers have decreasing priority. While the previous approaches were specifically developed for IP multicast, AMML is based on congestion control mechanisms used in ATM networks. In the rate-based method, the sender receives feedback explicitly in the form of the desired transmission rate for each layer. The sender initiates the feedback process by multicasting a "forward feedback packet." At each intermediate node, the ERICA algorithm [23] is used to calculate fair share of link bandwidth of the connection, and this is entered in the explicit rate (R_E) field of the forward feedback packet. When the packet reaches a receiver, the R_E field indicates the transmission rate the receiver's connection can support. Based on this information, receivers send "backward feedback packets" to the sender requesting specific transmission rates. Backward feedback packets are merged at intermediate nodes, concatenating the rate fields, and eliminating some if required according to specified criteria, so that the number of requested rates in a feedback packet does not exceed the number of layers the sender can support.

In the credit-based method, congestion feedback, as well as information about the number of receivers fully and partially receiving each video layer, propagates hop-by-hop back to the sender. The underlying principle is that an upstream node can send a certain number of packets to a downstream neighbor only if it has received an equal number of credits from the downstream node. The feedback packet eventually arriving at the sender indicates the total number of receivers fully and partially receiving each layer. The sender uses this information, as well as its buffer occupancy, to decide the number of video layers and the transmission rate of each layer.

4. Transcoder-Based Adaptation

An alternative approach to layered encoding and transmission is to use *video* (or *multimedia*) *gateways* at appropriate locations in the network to convert through transcoding a high bandwidth transmission into a transmission with appropriate bandwidth to accommodate groups of poorly connected receivers. In addition to configuring a session appropriately with gateways during start-up, receivers may be allowed to adapt to network congestion by dynamically identifying and requesting service from a node with better reception to serve as the gateway. Alternatively, or in addition, the gateway may use an adaptive rate-control algorithm to adjust its transmission in response to receiver feedback.

The two main considerations in developing a transcoder-based adaptation scheme are the design of the transcoding algorithm, and the placement or selection of the gateway to perform the transcoding.

In [24], Amir, McCanne and Zhang propose the following underlying model. The input format is converted into an intermediate representation by a decoder. This representation is transformed and delivered to the encoder, which produces a new bit stream in a new format (Fig. 4).

Instead of allowing a single intermediate representation, multiple intermediate formats are supported by the proposed transcoding model, allowing flexibility in choosing an encoder/decoder pair, and optimizing performance by enabling a higher level of intermediate

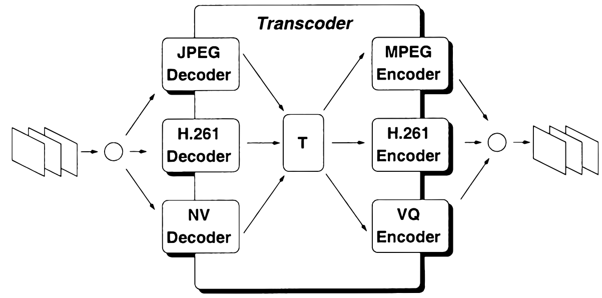

Fig. 4 Block diagram of transcoder proposed by Amir, et al. [24].

representation (such as DCT coefficients) to be used instead of decomposing the input stream into pixel format. Using a selectable intermediate format is also meant to give more flexibility in transformations to achieve a target output rate, including temporal and spatial decimation and/or frame geometry conversion. The transcoder is configured by an external control interface through which parameters such as encoding and decoding formats, output rate, compression parameters, etc. can be specified. A more flexible scheme for configuration and control of transcoders has been suggested in [25].

Kouvelas, Hardman and Crowcroft [26] present a control scheme that automatically configures transcoders within the multicast tree to support branches with bad reception. A group of receivers affected by a bottleneck tries to locate an upstream receiver with better reception to provide a customized, transcoded version of the session stream by multicasting request messages. To prevent requests from multiple receivers in the group from proliferating, a requester delays its request by an interval proportional to its distance from the stream source plus a small random interval. If the requester receives an identical request during this delay, it cancels its own request.

5. Discussion

5.1 Signaling Mechanism

A primary issue in rate adaptation schemes is the choice of the indicator which signals congestion in the network. The feedback signal may be explicit, or network assisted — for example, occupancy of a buffer along the transmission path, request for a specific transmission rate from the network or receivers, or source quench messages. The feedback signal may also be implicit, or directly from receiver to sender without network intervention — for example, packet loss, packet delay, delay jitter, throughput, or CSMA/CD collisions. The main criteria in evaluating the signaling mechanism are the promptness and reliability with which congestion is indicated.

A majority of the schemes discussed in this study use packet loss as an indicator (the loss based schemes by definition, RLM [18] and ALT [22] among the receiver based schemes, and the transcoder based scheme of Kouvelas [26], et al.). An important consideration in this case is the time interval over which loss rate is measured. If the time interval varies, for example, as a scaling mechanism, loss arising from a transient congestion may appear as a high loss rate when the measurement interval is relatively short, and may appear insiginificant over a long measurement interval.

Since many existing multimedia protocols and tools in use over the Internet are based on RTP, the use of RTCP reports as feedback in packet loss-driven control schemes eliminates the need for an additional feedback mechanism. Of the schemes discussed in this paper, LDA [21], ALT [22], and the scheme proposed by Busse, et al. [15] specify the use of RTCP receiver reports as feedback.The RTCP scaling mechanism increases the interval between receiver reports as the size of the multicast group increases, to keep the RTCP control traffic at a constant level. Assuming that the sender makes rate adjustments at fixed intervals, this implies that a decreasing fraction of receiver responses is sampled as the number of receivers increases. A similar sampling of the receiver population is achieved by a probabilistic polling method in the scheme proposed by Bolot, et al. [16]. This is discussed in the next section.

In the LDA algorithm [21], in addition to the loss feedback which drives the control loop, the sender considers the round-trip delay computed from RTCP receiver reports, and also forms an estimate of the bottleneck bandwidth for each receiver using an additional signaling mechanism based on the packet pair approach described by Bolot [17]. The additional information is used to determine the additive rate increase parameter array AIR_i, so as to limit the rate of increase to that of of an equivalent TCP connection.

In the ThinStreams protocol [19], the receiver estimates the expected throughput (on the basis of the number of fixed bandwidth thin streams it subscribes to) and compares it to the actual received throughput in order to measure congestion.

The buffer based adaptation schemes based on sender adaptation and the credit-based feedback mechanism considered in AMML [28] use the occupancy of a buffer on the transmission path as a measure of congestion. In the rate-based scheme of AMML, the sender obtains a set of cumulative transmission rates as feedback, based on which it determines the output rate of each encoded layer. In general, all the explicit signaling mechanisms discussed need router or switch support and impose an additional overhead in the form of a special hop-by-hop feedback mechanism.

5.2 Multicast Scalability of Control Mechanism

Sender-based schemes

The schemes based on buffer occupancy discussed here are intended solely for unicast applications, and scalability considerations do not apply to them. In loss based adaptation schemes, the scalability of the protocol is likely to be determined by the scaling properties of the feedback mechanism by which receivers report loss to the sender.

In the LDA scheme [21] and the scheme proposed by Busse, et al. [15], RTCP receiver reports convey feedback information back to the sender. The frequency of the feedback is governed by the RTCP scaling mechanism, which adjusts the interval between receiver reports so that the RTCP control traffic makes up no more than 5% of the total traffic. In case of a long feedback interval, it is possible that the reported congestion may have cleared by the time the sender reacts. The IETF is considering modifying RTCP to report losses seen in a small interval prior to the report, instead of the average loss between the sending of two consecutive reports.

The feedback process in the scheme proposed by Bolot, et al. [16] relies on a probabilistic polling mechanism. Probe messages are multicast in order to elicit congestion reports from receivers. The messages are sent in successive rounds, and the number of receivers eligible to respond on a particular round is restricted by requiring them to match a randomly generated key. The sender first transmits SIZESOLICITED probe messages to which any receiver with a matching key is eligible to respond. The length of the key is gradually reduced in successive rounds, until a response is received. A logarithmic relationship is shown to exist between the number of receivers n and the average probing round in which a receiver response is first received.

Subsequently, the sender transmits probe messages to which only receivers experiencing congestion respond upon matching the random key. The fraction of congested receivers in the group is estimated from the number of elapsed rounds between a response to a SIZESOLICITED probe, and the first response reporting congestion. This method has the advantage that the maximum discovery time of a congested receiver is independent of the number of receivers. In the worst case, with no congested receivers, the congestion discovery time is bounded by an interval equal to $2 * l * rtt_{max}$ where rtt_{max} represents the worst case round trip time for a probe message and l is the length of the random key (a length of 16 is seen to be adequate for handling up to 10000 receivers), and it decreases logarithmically from the worst case value as the number of receivers increases.

Receiver-based schemes without rate adaptation

In receiver based, layered transmission schemes that do not employ rate adaptation by the sender, the scalability is determined by the join experiments conducted by receivers trying to obtain higher QoS by adding an enhancement layer. A failed experiment results in a transient increase in congestion before the receiver learns of its failure and drops the layer. In the absence of a scaling mechanism, the frequency of these transients, as also the possibility of join experiments interfering with each other become the limiting factors.

Scalability can be achieved by increasing the minimum interval between join experiments in proportion to the overall group size. This comes at the cost of an increased convergence time to reach a stable subscription level. In RLM [18], scalability is realized by a *shared learning* mechanism, in which a receiver broadcasts a join experiment announcement before performing the join experiment. If a receiver waiting to perform a join experiment to the same layer experiences congestion during the announced experiment, it deduces failure of the join experiment without performing the experiment itself, and backs off the associated join-timer. In this way, the number of failed join experiments does not increase in proportion to the size of the group. However, a receiver cannot similarly learn from the success of a join experiment because a bottleneck may be present on its delivery branch, which was unaffected by the experiment. Also, the possibility of join experiments of different multicast sessions interfering with each other (due to accidental synchronization of the experiments) is not addressed.

In ThinStreams [19], the start of join-experiments within a session are synchronized — hence, receivers conduct join experiments simultaneously. This is done with the aim of minimizing the frequency of join experiments which are likely to result in transient congestion if they fail. However, this appears likely to limit the scalability of the join mechanism — if a large number of receivers all conduct join-experiments to different layers, the joint failure of all of these experiments is an increasingly unreliable indicator that each of those experiments would have failed if conducted separately. The issue of interference between join experiments in different sessions is addressed by randomizing the join start times of different sessions.

Receiver-based schemes with rate adjustment

In general, similar scalability considerations should apply to joins in receiver based schemes with rate adjustment. No coordination mechanisms are discussed for joins and leaves in either ALT [22] or AMML [28], and in their absence (that is, if receivers add and drop layers completely independently of each other), the join mechanism may scale poorly. It is not obvious how the frequency of joins would be affected by incorporating sender adaptation into a multi-layered scheme. It is possible that due to dynamic changes in the output rates of the layers, receivers may in fact add or drop layers more frequently than in the absence of sender adaptation.

The ALT scheme uses RTCP receiver reports to drive adaptation of each layer by the sender, and the same trade-offs discussed for sender-driven schemes apply in this case. In AMML, both rate-based and credit-based mechanisms use hop-by-hop feedback messaging originating from the receivers. At intermediate nodes, state information from feedback packets which arrive close together are merged, and a new feedback packet is generated and sent upstream. Consequently, the receiver feedback mechanism should scale better in this case.

Transcoder-based schemes

In the transcoder based scheme of [26], during the process of identifying a transcoder for a group of bottlenecked receivers, the number of request messages is scaled by scheduling a message for an interval proportional to the distance of the originator from the source instead of sending it immediately, and cancelling it if a duplicate request is received during the scheduled interval. A similar scaling mechanism is followed for response messages to requests by candidate transcoders.

5.3 Fairness and Interaction with TCP

TCP-like congestion control

In the buffer based adaptation scheme of Jacobs, et al. [10], as well as the transcoder based scheme of Kouvelas, et al. [26], the TCP congestion window is used to drive the rate adaptation process, without the mandatory retransmission of lost packets done in TCP. The buffer based scheme is shown to result in a roughly equitable distribution of bandwidth between TCP-based and real-time traffic. As discussed in the introduction, however, the use of TCP congestion control in real-time multimedia applications has a number of disadvantages, although the effect of the abrupt rate-reduction in TCP may be minimized by the use of local buffering, at the cost of an additional delay.

The LDA scheme [21] borrows a number of features from TCP congestion control in order to improve fairness to competing connections, although it responds to loss notifications in RTCP reports instead of relying on per-packet acknowledgments. In LDA, the additive rate of increase (AIR_i) is adjusted according to the estimated bottleneck bandwidth: $\mathrm{AIR}_i = \mathrm{AIR}(1 - r/b)$, where r is the current transmission rate, b is the estimated bottleneck bandwidth, and AIR is the current additive rate of increase. Additionally, AIR_i is maximally limited to the average increase r_{incr} in the rate of a loss-free TCP connection over the rate-adaptation period used in the LDA algorithm. Finally, if the RTCP

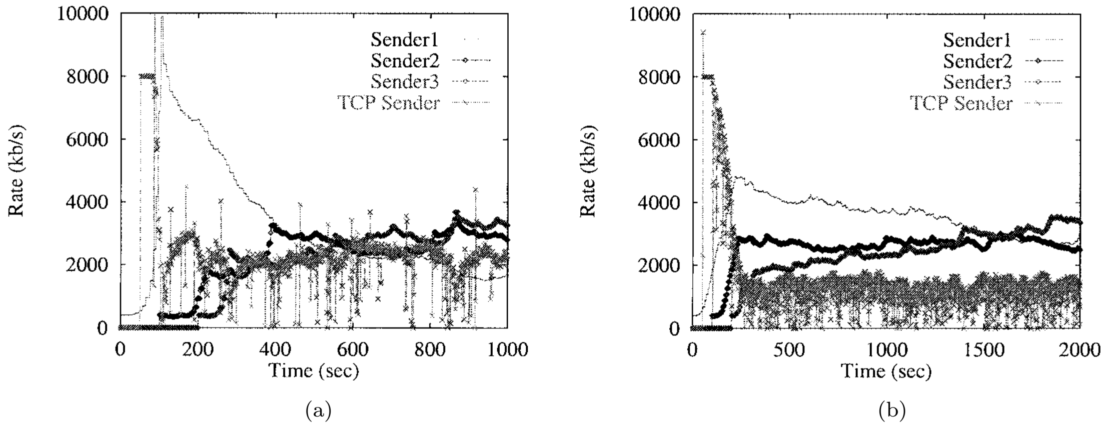

Fig. 5 Sisalem, et al. [21]: Bandwidth of competing LDA and TCP connections with available bandwidth = 10 Mbps, and end-to-end delays 1 ms (a) and 100 ms (b). Bandwidth is evenly distributed in (a), but the TCP connection receives a smaller share than the LDA connections in (b).

feedback reports multiple lost packets, the rate is reduced by a factor proportional to the number of lost packets, in order to achieve TCP like behavior. As shown in Fig. 5, fairness to a competing TCP connection may be limited by the resolution with which loss can be reported by the feedback mechanism. Fairness to competing LDA connections is also demonstrated in Fig. 5. As with TCP connections, a smaller bandwidth is allocated to connections having long round trip delays and to connections traversing a larger number of hops.

Other loss-based schemes

The other loss-based schemes [15], [16] follow the rate adaptation model of Fig. 2. Instead of responding to each loss notification by scaling down the rate as in TCP, the rate is scaled down only when the loss percentage exceeds a certain threshold — this may be more appropriate for real-time applications since they can typically absorb a small amount of loss without significant degradation. Also, the reduction factor is constant, independent of the number of lost packets. The degree of fairness of the algorithms depends on the specific parameter values used, and equitable distribution of bandwidth among connections is not guaranteed. Unfairness to TCP connections increases as the loss threshold (λ_c) is increased.

Bandwidth distribution between peer connections may also be unfairly skewed if, at the onset of congestion, one connection is attempting to transmit at a higher rate than the other. Assuming both connections transmit under similar conditions, they would both reduce their rates by similar factors until congestion is removed, resulting in the first connection receiving a higher share of the bandwidth. More equitable distribution may be achieved by making the rate reduction factor dependent on the amount of loss, as in LDA. Also, if all the competing connections use the same value of loss threshold, connections with larger number of hops (and hence experiencing more lost packets) would receive a smaller share of bandwidth, as in TCP.

Receiver-based schemes

In RLM [18], a receiver adapts its received signal by adding or dropping a layer, and thus changes its received bandwidth with the granularity of one encoded layer. For typical layered encoding schemes, this can result in unfairness to competing connections due to large changes in bandwidth arising from adaptive action. A receiver with excess capacity on its delivery path might capture the bulk of this bandwidth by subscribing to an additional layer, instead of increasing its share in small increments, as is done in TCP and in a number of the adaptive schemes discussed here. The potential for this happening also exists in ALT [22] and credit based AMML [28]. In the ThinStreams protocol, this problem is addressed by splitting up each encoded layer (thick stream) into several smaller bandwidth layers (thin streams), so that bandwidth is added and dropped in smaller increments.

In ALT, the rate adaptation of each layer is stated to follow the model of Fig. 2, but details are not discussed. In the rate based scheme proposed under AMML, the ERICA algorithm is used to calculate the connection's fair share of bandwidth at each router, and the bandwidth delivered to a receiver is constrained by this fair share. When competing connections all employ AMML and have the same propagation delay and video rate, both rate based and credit-based algorithms are shown to distribute bandwidth fairly among competing connections.

Among competing connections employing the ThinStreams protocol, fair sharing of link bandwidth is supported by scaling the *leave_threshold* (loss threshold at which a receiver drops its highest layer) as an exponentially decreasing function of the number of groups (G) the receiver has joined, leave_threshold = $G * R * e^{(1-G)/8}$. R represents the bandwidth of a thin stream. A multicast session in which receivers subscribe to more groups (and hence receive higher bandwidth) has a lower leave-threshold, and its receivers drop groups faster on experiencing congestion than receivers in a group with lower QoS.

In RLM and in ThinStreams [19], receivers start with the lowest level of QoS upon joining a session, and both protocols have mechanisms to allow new receivers to converge quickly to their stable subscription levels by conducting join experiments more frequently. In RLM, a receiver wanting to join a lower level than the level of an on-going join experiment is allowed to perform its own join experiment simultaneously. In ThinStreams, the *hold-off-time* timer which triggers a join experiment increases proportionally with the number of groups a receiver is subscribed to.

5.4 Loss and Bandwidth Utilization in a Heterogeneous Network

Sender-based schemes

In a multicast session, the sender has the problem of determining a single optimal transmission rate in response to feedback from receivers which may be heterogeneous in their computational ability, network connectivity, and need for transmission quality, and may report very different levels of loss depending on their connectivity. One possible approach is to adjust the transmission rate according to the most poorly connected receiver. This approach is taken in LDA [21]. Another possible approach is to allow a certain fraction of the total number of receivers to report congestion before entering the congested state, and reducing the transmission rate accordingly. This approach is followed by Bolot, et al. [16]. Both approaches are examined by Busse, et al. [15]. The first approach may result in the majority of participants receiving low quality transmissions because of one poorly connected receiver. The second approach may result in a certain number of receivers suffering from poor or unacceptable quality due to high packet loss on their delivery paths. In either case, a highly heterogeneous environment will result in poor bandwidth utilization along some delivery paths.

Receiver-based schemes

Sender-based adaptation is fundamentally limited in its ability to accommodate a heterogeneous group of receivers, because of the need to adapt a single transmission to meet the needs and capabilities of different members of the group. Receiver-based, layered adaptation methods are inherently better equipped to handle receiver heterogeneity. They allow each receiver to tailor the signal it receives by adding or dropping enhancement layers, by subscribing to or leaving multicast groups. This approach also allows higher utilization along all delivery paths.

In RLM [18], the video codecs determine the bandwidth of each layer, and QoS adjustments by the receiver have the granularity of one encoded layer. As a result, RLM may not be able to utilize available bandwidth as efficiently as the other layered adaptation schemes, which refine the same basic approach in different ways.

In ThinStreams [19], bandwidth utilization is improved by allowing the receiver to adapt its delivery by adding or dropping lower bandwidth thin streams, instead of the original encoded layers, or thick streams. However, this advantage is contingent upon the decoder being able to re-assemble partial thick streams, and hence upon the particular encoding scheme used. If partial thick streams cannot be re-assembled, the bandwidth used to transmit the corresponding thin streams is wasted.

The ALT scheme [22] enables the sender to adapt the number of layers and the rate of each layer in response to receiver loss feedback, thus allowing bandwidth adaptation in smaller increments.

In rate-based AMML [28], the sender adapts the transmission rate of each layer based on the available "fair-share" bandwidth along each delivery path. In credit-based AMML, the transmission rates are controlled by a hop-by-hop flow-control mechanism. For a small multicast session, both approaches are shown to result in nearly 100% utilization, similar to other ABR schemes.

Transcoder-based schemes

Transcoder-based schemes have the capability of delivering appropriate QoS levels to a heterogeneous group of receivers without incurring the overheads associated with layered encoding and transmission. In a session with a static configuration of transcoders, and under the assumption that each transcoder serves a homogeneous cluster of receivers, each transcoder can serve its group with very high utilization. In a dynamic configuration as envisaged in [26], however, there may be one or more receivers behind the bottleneck link which are still subscribing to the original stream instead of the transcoded stream. This makes a "slow-start" phase necessary when the transcoder initiates transmission to avoid further congestion in the bottleneck link, and this will temporarily result in sub-optimal utilization.

5.5 Implementation Cost and Complexity

In general, adaptation schemes that do not require network intervention are likely to be easier to implement than protocols which incorporate hop-by-hop mechanisms and require network intervention. Among the adaptation schemes discussed here, the following are network assisted.

- In the buffer based adaptation scheme of Kanakia, et al. [9], the sender requires state information about the bottleneck router. The state information propagates in either the forward or reverse direction, and each router along the transmission path either updates the state information, or passes it unmodified.

- In AMML [28], intermediate nodes are responsible for periodically collecting feedback messages from downstream routers, and merging the state information into a new feedback packet which is sent upstream. Additionally, in rate-based AMML, the sender multicasts a feedforward message; each switch implements a fairness algorithm (ERICA) to calculate bandwidth share for the connection, and updates state information in the feedforward message accordingly. In credit-based AMML, the switch maintains state information about the number of packets served at output links, and sends "credits" to the upstream router in the merged feedback message.

Layered encoding schemes require complex encoding and decoding systems. The additional complexity makes it desirable to allow feedback from the receivers (for example, as in ALT and AMML), so that the creation of layers reflects the receiver interest, and does not result in layers with few subscribers, or poorly received layers. The need to synchronize decoded streams at the receiver also adds to the end-to-end delay. The ThinStreams algorithm [19] introduces a large additional overhead at the receiver arising from the need to decode and synchronize a large number of layers. Layered transmission systems with rate adjustment have additional complexity in the encoding system, since the number of layers and the output rate of each layer must be adjusted in response to feedback. In particular, in ALT [22], the rate adaptation process requires per-receiver state information to be maintained by the sender, and this is likely to become significant in a large multicast session. Also, it is not clear if any of the existing encoding system permit adjustment of the output rate of individual layers, without modification.

Transcoder-based systems require the implementation of transcoding systems (multiple formats and rate adaptation for congestion control may have to be supported) at gateway nodes.

5.6 Responsiveness and Perceived Quality

The adaptive scheme should detect and react to congestion quickly, to minimize poor reception quality and interruptions at the receiver. At the same time, the adaptation should not result in abrupt changes or oscillations in the reception quality.

The adaptive schemes employing TCP-like congestion control react to every packet loss by halving the rate. This can result in rapid changes in the transmission rate (for example, Fig. 6) which may be perceived as unpleasantly abrupt changes in quality at the receiver. In the LDA scheme [21], a TCP-like congestion control mechanism is used in which the multiplicative rate reduction factor can be reduced to obtain a smoother change in the rate at the cost of a longer convergence time, but without sacrificing the goal of fairness to competing TCP connections. In the buffer-based scheme of Jacobs, et al. [10], abrupt changes may be smoothed out by the use of local buffering, but this is reported to add a delay of a few seconds, undesirable in real-time applications. The scaling property of the adaptive scheme is critical for timely response in a large multicast group. Among adaptive schemes with rate adaptation by the sender, in schemes using RTCP feedback the worst-case congestion discovery time is determined by the RTCP feedback interval, which increases proportionally with the number of receivers. The probabilistic polling proposed by Bolot, et al. [16] results in a worst-case congestion discovery time independent of the number of receivers, and proportional to the maximum round-trip time. The rate based algorithm used in AMML can reply to congestion status by taking one round trip time from bottleneck link to the source. Credit based algorithm take longer time to converge due to its use of incremental rate changes in response to network feedback. Again, the quick response is at the cost of abrupt changes in peceivetive quality.

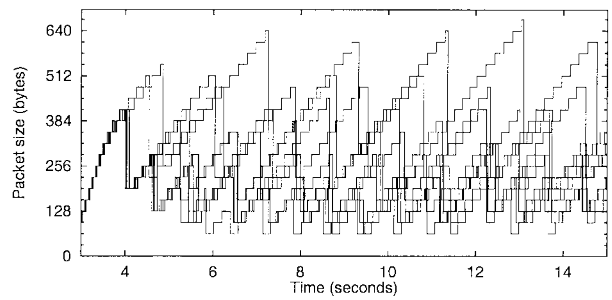

Fig. 6 Adaptation of packet size in the TCP-like congestion control scheme used by Kouvelas, et al. [26].

The buffer based scheme of Kanakia, et al. [9] tries to compensate for the latency of the reported buffer occupancy by attempting to model the evolution of the system state in order to obtain more up-to-date feedback. However, the impact of the predictive mechanism is not studied independently in reported experiments. In general, schemes using explicit signaling mechanisms may be expected to react pro-actively to prevent congestion and significant packet loss.

The responsiveness of layered transmission schemes may be affected by the join and leave latencies of the underlying multicast protocol. For example, the default leave latency of IGMP [33] is 2 seconds. Fluctuations and abrupt changes in quality may occur in RLM [18], due to the large granularity with which rate adaptation takes place. On the other hand, the small adaptation increments in ThinStreams [19] may result in an overly slow response, since the latency in adding or dropping each layer depends on the underlying multicast protocol. It may be preferable to allow the receiver to drop a certain number of layers in case of an abrupt reduction in throughput, or occurrence of packet loss. In general, the constraints imposed by a multi-layered encoding scheme are likely to result in a transmission of inferior quality compared to a single encoded layer with the equivalent transmission bandwidth.

The transcoding process in transcoder-based schemes may introduce a significant delay. In [26], dynamic configuration of transcoders within the session is envisioned in response to congestion. This requires a transcoding initiation process to identify the transcoder, and coordinate the switching of a group of bottlenecked receivers to the transcoded stream. Simulations by the authors show that the transcoding initiation process occupies several seconds.

6. Conclusion

We have presented a review of some recent work on rate-adaptive control schemes for multimedia applications. The adaptation schemes fall into three broad categories — sender-driven, receiver-driven, and transcoder-based. The sender-driven schemes discussed here adapt the sender rate based on either buffer occupancy, or on loss rate feedback from the receivers. The receiver-driven schemes discussed all use multi-layered encoding and transmission systems in order to allow a receiver to select a service with appropriate QoS. Some of them combine this approach with sender adaptation of the number of encoded layers and output rate of each layer based on feedback. Transcoder-based schemes rely on transcoders to transform the original high bandwidth source stream into a stream with appropriate bandwidth to serve poorly connected or congested receivers.

The loss-based, sender-driven schemes discussed here have relatively low overhead and simple implementations. In comparison, layered transmission schemes add complexity and delay to the encoding and decoding systems, and transcoder-based systems require implementation of transcoding systems at routers, and have possible concerns about increase delay and security. At the same time, sender-driven schemes are limited in their ability to accommodate a heterogeneous group of receivers which differ in their connectivity or the amount of congestion on their delivery paths. They are better-suited to multicast sessions with a homogeneous group of receivers, distributed over a relatively small area. In larger, heterogeneous sessions, transcoder-based schemes and layered transmission schemes are better suited. Transcoder-based schemes may be preferred in a session which has diverse receivers with different connectivities. In an environment with dynamically changing congestion, layered transmission schemes may be preferable to avoid having to dynamically locate and configure transcoders.

A number of adaptation schemes assume RTP-based communication, and most of them use RTCP receiver reports to convey feedback to the sender. This approach has the convenience of avoiding a separate feedback mechanism. Bolot, et al. [16] instead use a probabilistic polling scheme to estimate receiver congestion, and show that this results in a worst-case congestion discovery time independent of the number of receivers. A separate messaging scheme is also proposed in AMML [28], but it requires significant router intervention. AMML also assumes the availability of prioritized service at routers.

The congestion-control properties of the rate-adaptive schemes are of importance in multimedia communication over the Internet. A number of schemes use TCP-like congestion control schemes, and demonstrate fairness to competing connections, including TCP connections. However, strictly TCP-like congestion control may result in sharp reductions in the transmission rate, and possibly unpleasant reception quality. In the LDA adaptation scheme, by using an adjustable rate-reduction factor, a smoother adjustment is obtained at the cost of slower convergence. In general, the effect on the perceived quality of the trade-off between responsiveness and a smooth adjustment of the rate is not addressed in detail, and appears to be an important area for further investigation.

References

[1] H. Schulzrinne, "Transport protocols for multimedia," Tutorial Infocom'98.

[2] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource ReSerVationProtocol (RSVP)—Version 1 Functional Specification," Internet Engineering Task Force, Internet Draft, June 1997.

[3] P. Pan and H. Schulzrinne, "YESSIR: A simple reservation mechanism for the Internet," in International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV97), Cambridge, England, July 1998.

[4] S. Jamin, S.J. Shenker, and P.B. Danzig, "Comparison of measurement-based admission control algorithms for controlled-load service," Proc. IEEE INFOCOM '97, April 1997.

[5] H. Zhang and S. Keshav, "Comparison of rate-based service disciplines," Proc. ACM SIGCOMM'91, Zurich, Switzerland, Sept. 1991.

[6] K. Nichols, V. Jacobson, and L. Zhang, "A two-bit differentiated services architecture for the Internet," Internet Draft, Internet Engineering Task Force, Nov. 1997.

[7] T. Sakatani, "Congestion avoidance for video over IP networks," Proc. First International COST 237 Workshop on Multimedia Transport and Teleservices, Vienna, Austria, pp.256–273, Springer, Nov. 1994.

[8] K. Jeffay, D.L. Stone, T. Talley, and F.D. Smith, "Adaptive, best-effort, delivery of audio and video data across

packet-switched networks," Intl. Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV92), Nov. 1992.

[9] H. Kanakia, P. Mishra, and A. Reibman, "An adaptive congestion control scheme for real-time packet video transport," SIGCOMM Symposium on Communications Architectures and Protocols, San Francisco, California, pp.20–31. ACM/IEEE, Sept. 1993.

[10] S. Jacobs and A. Eleftheriadis, "Streaming video using dynamic rate shaping and TCP congestion control," Journal of Visual Communication and Image Representation, vol.9, iss.3, pp.211–222, Sept. 1998.

[11] S. Jacobs and A. Eleftheriadis, "Real-time dynamic shaping and control for Internet video applications," Workshop on Multimedia Signal Processing, Princeton, NJ, June 23-25, 1997.

[12] A. Eleftheriadis, S. Pejhan, and D. Anastassiou, "Constrained and general dynamic rate shaping of compressed digital video," Proc. 2nd IEEE International Conference on Image Processing, Washington, DC, Oct. 1995.

[13] Z. Chen, S.M. Tan, R.H. Campbell, and Y.Li, "Real time video and audio in the World Wide Web," World Wide Web Journal, Vol.1, Jan. 1996.

[14] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: A transport protocol for real-time applications," RFC 1889, Jan. 1996.

[15] I. Busse, B. Deffner, and H. Schulzrinne, "Dynamic QoS control of multimedia applications based on RTP," Computer Communications, vol.19, no.1, pp.49–58, Jan. 1996.

[16] J.-C. Bolot, T. Turletti, and I. Wakeman, "Scalable feedback control for multicast video distribution in the Internet," in SIGCOMM Symposium on Communications Architectures and Protocols, London, England, pp.58–67, ACM/IEEE, Aug. 1994.

[17] J.-C. Bolot, "End-to-End packet delay and loss behavior in the Internet," in SIGCOMM Symposium on Communications Architectures and Protocols, London, England, pp.58–67, ACM/IEEE, Aug. 1994.

[18] S. McCanne, V. Jacobson, and M.Vetterli, "Receiver-driven Layered Multicast," in SIGCOMM Symposium on Communications Architectures and Protocols, Stanford, CA, Aug. 1996.

[19] L. Wu, R. Sharma, and B. Smith, "Thin streams: An architecture for multicasting layered video," the 7th International Workshop on Netwwork and Operating Systems Support for Digital Audio and Video (NOSSDAV97), St. Louis, USA, May 1997.

[20] E. Amir, S. McCanne, and R. Katz, "Receiver-driven bandwidth adaptation for light-weight sessions," ACM Multimedia'97, Seattle, Nov. 1997.

[21] D. Sisalem and H. Schulzrinne, "The loss-delay adjustment algorithm: A TCP-friendly adaptation scheme," Proc. International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV98), Cambridge, England, July 1998.

[22] D. Sisalem and F. Emanuel, "QoS control using adaptive layered data transmission," Proc. IEEE Multimedia Systems, Austin, Texas, Dec. 3, 1997, June 1998.

[23] R. Jain, S. Kalyanaraman, R. Goyal, S. Fahmy, and R. Viswannathan, "ERICA switch algorithm: A complete description," ATM Forum/96-1172, Aug. 1996.

[24] E. Amir, S. McCanne, and H. Zhang, "An application level video gateway," Proc. ACM Multimedia '95, San Francisco, Nov. 1995.

[25] E. Amir, S. McCanne, and R. Katz, "An active service framework and its application to real-time multimedia transcoding," in SIGCOMM Symposium on Communications Architectures and Protocols, Vancouver, BC, Canada, Sept. 1998.

[26] I. Kouvelas, V. Hardman, and J. Crowcroft, "Network adaptive continuous-media applications through self organized transcoding," Proc. Network and Operating Systems Support for Digital Audio and Video (NOSSDAV98), Cambridge, UK, July 8–10, 1998.

[27] S. Floyd, V. Jacobson, S. Liu, S. McCanne, and L. Zhang, "A reliable multicast framework," IEEE/ACM Trans. Networking, vol.5, no.6, pp.784–803, Dec. 1997.

[28] B.J. Vickers, C. Albuquerque, and T. Suda, "Adaptive multicast of multi-layered video: Rate-based and credit-based approaches," Proc. IEEE Infocom '98, San Francisco, 1998.

[29] The Flow Consortium c/o Ascom Nexicon Inc., Quantum Flow Control, Version 2.0, July 1995.

[30] K. Brandenburg and T. Sporer, "NMR and masking flag: Evaluation of quality using perceptual criteria," Proc. AES 11th Intl. Conf, pp.169–176, 1992.

[31] J.G. Beerends and J.A. Stemerdink, "A perceptual audio quality measure based on a psychoacoustic sound representation," J. Audio Eng. Society, vol.40, no.12, Dec. 1992.

[32] O. Verscheure, P. Frossard, and M. Hamdi, "MPEG-2 video serivices over packet networks: Joint effect of encoding rate and data loss on user-oriented QoS," The 8th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV98), Cambridge, England, pp.257–164, July 1998.

[33] W. Fenner, "Internet Group Management Protocol, Version 2," RFC 2236, Nov. 1997.

Xin Wang received her B.S. and M.S. degree in telecommunications engineering from Beijing University of Posts and Telecommunications, Beijing, China, in 1990 and 1993 respectively. She is currently working toward her Ph.D. in the Department of Electrical Engineering at Columbia University, New York, USA. Her current research interests include QoS provisioning and pricing for the Internet, multimedia adaptive services, and multicast routing and reliability. She is a student member of IEEE.

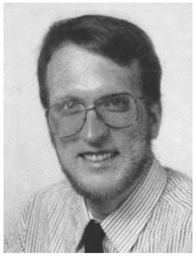

Henning Schulzrinne received his undergraduate degree in economics and electrical engineering from the Technische Hochschule in Darmstadt, Germany, in 1984, his MSEE degree as a Fulbright scholar from the University of Cincinnati, Ohio and his Ph.D. degree from the University of Massachusetts in Amherst, Massachusetts in 1987 and 1992, respectively. From 1992 to 1994, he was a member of technical staff at AT\&T Bell Laboratories, Murray Hill. From 1994–1996, he was associate department head at GMD-Fokus (Berlin), before joining the Computer Science and Electrical Engineering departments at Columbia University, New York. His research interests encompass real-time, multimedia network services in the Internet and modeling and performance evaluation. He is an editor of the Journal of Communications and Networks and IEEE Communications Society editor of the IEEE Internet Computing Magazine. He co-chairs the IEEE Communications Society Internet Technical Committee and is vice chair of the IEEE Communications Society Technial Committee on Computer Communications. He has been vice general chair of IEEE Infocom and will be co-technical chair of that conference in 2000. Protocols co-developed by him are now Internet standars, used by almost all Internet telephony and multimedia applications. He is co-author of the Real-Time Protocol (RTP) for real-time Internet services,the signaling protocol for Internet multimedia confernces and telephony (SIP) and the stream control protocol for Internet media-on-demand (RTSP). He currently serves as feature editor of IEEE Communications Magazine and IEEE Communications Society editor for the IEEE Internet Computing Magazine.